==

BAA 05-43 PROPOSER INFORMATION PAMPHLET

==

The Defense Advanced Research Projects Agency (DARPA) often selects its research efforts through the Broad Agency Announcement (BAA) process. The BAA will be posted directly to FedBizOpps.gov, the single government point-of-entry (GPE) for Federal government procurement opportunities over $25,000. The following information is for those wishing to respond to the Broad Agency Announcement.

Integrated Learning, SOL BAA 05-43, Proposals Due: Initial Closing: 14 September 2005, Final Closing: 11 July 2006, POC: Dr. Tom Wagner, DARPA/IPTO; FAX: (703) 741-7804

PROGRAM OBJECTIVES AND DESCRIPTION

Introduction

The Defense Advanced Research Projects Agency (DARPA) Information Processing Technology Office (IPTO) is soliciting proposals for a new program called Integrated Learning. This program will develop computer software, called an Integrated Learner, which learns general plans or processes from human users by being shown one example. Integrated Learners will accomplish this by opportunistically assembling knowledge from many different sources, including generating it by reasoning, in order to learn. Applications include learning air-tasking-order (ATO) planning or CAD design processes. The Integrated Learning technology will enable low-cost military decision/planning support systems.

Each individual proposal must specify a complete Integrated Learner solution, including but not limited to (1) an Integrated Learning software framework (e.g., common knowledge structures, integration software), (2) an Integrated Learning Problem Solver, (3) reasoning components that will be included in the Integrated Learner (e.g., domain planners), (4) simulation components that will be included in the Integrated Learner (i.e., tools that enable the learner to ask what-if questions about the world), (5) general world knowledge the learner will use, (6) domain knowledge the learner will use, (7) a cyberspace application, (8) a specification of the hierarchical plan or task models the learner will learn and produce, (9) the means of interacting with the human user, and (10) evaluation plans and metrics. More information on these components follows below.

The program is expected to have four 12-month phases. Only Phase I will be funded initially, however, proposers should address all four phases. DARPA will host an Industry Day for the Integrated Learning program on August 3, 2005. For more details and registration information please go to https://www.schafertmd.com/intlearning2005. Additional BAA details follow.

Program Goals

This program has four goals:

(1) To create Integrated Learners, a new computer software technology in which the software assembles knowledge from many different sources, including generating it by reasoning, in order to learn.
(2) To apply Integrated Learning software to cyberspace planning domains in which the software must learn complex hierarchical task or process models.
(3) To evaluate the Integrated Learning software's efficacy at learning these models.
(4) To transition the Integrated Learning software for military application and/or to create a new technology base from which future projects can draw to accomplish their goals.

Technical Overview

The Integrated Learning program will create a new kind of learning system in which learning is an *integrated problem solving process* where the learner opportunistically assembles knowledge from many different sources, including generating it by reasoning, in order to learn. The challenge problem for the learner is to learn a complex task model or generalized plan by being shown how to perform some task <u>*only once.*</u> To accomplish this, the learner must combine the limited observational data with domain knowledge, world knowledge, reasoning, and simulation (asking what-if questions) in order to assemble the body of knowledge necessary to generate the models. Learners in this program will not be exposed to large numbers of training instances as a primary learning input mechanism.

The expression *learning as an integrated problem solving process* identifies two important ideas: (1) learners in this program will be integrated in a meaningful two-way fashion with other components in a cognitive system and able to utilize their knowledge and their reasoning in the learning process, and (2) learners will regard learning as a problem to solve rather than a rote series of steps or operations to perform. On the latter point, this means that Integrated Learners will (1) have explicit learning goals and formulate plans to achieve them, (2) keep track of what they don't know and what they need to know, (3) form hypothesis and track uncertainties associated with them, (4) be both opportunistic and process driven in their control, (5) assemble knowledge from multiple sources and build on that knowledge. This is not a complete or exhaustive list. From an executive level, Integrated Learners attempt to "figure things out" rather than execute a set of predefined/static algorithmic steps. This approach will yield learning systems that are more flexible; where the learners are able to utilize many different sources of information, process information in many different forms, and proactively work with reasoning and simulation components to generate desired information. Integrated Learners will also be more robust – tolerant of errors in information and tolerant against missing information because the learner can draw on whatever information or reasoning is available to support learning and can use multiple sources to corroborate or negate hypotheses.

An example Integrated Learner is shown in Figure 1. Note that the learner incorporates world knowledge, domain knowledge, several types of sophisticated reasoning and simulation, and a module for conventional statistical machine learning algorithms. These other components or modules are *tools* that the Integrated Learner employs during learning to generate knowledge that it needs to achieve its learning goals. This interconnected view is very different from the algorithmic focus of statistical machine learning algorithms where the algorithm has one "input pipe" through which training data (of a very specific form) flows. Integrated Learners have many “pipes” and must be able to manipulate many different forms of information and even trade off different types of information and reasoning. Integrated Learners can also interact directly with a human user to fill learning information needs. However, the learner must perform a cost/benefit trade-off analysis before invoking the human as the human interaction may be a more “expensive” option than other computational options the learner may have.

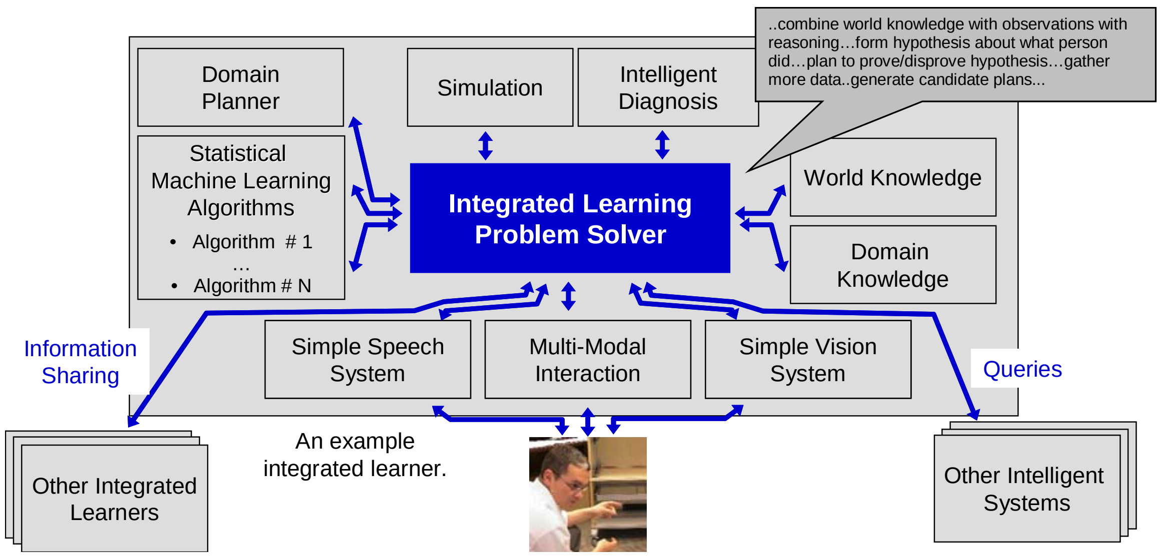

Figure 1 - An Example Integrated Learner for Physical Domains

The motivating challenge problem for Integrated Learning is to learn (symbolic) hierarchical task models, complex processes, or generalized plans by watching a human user perform a task *just once*. Figure 2 shows an example taken from a physical domain in which the learner must produce a generalized task model for assembling an object after being show the assembly process, once, by a human user.

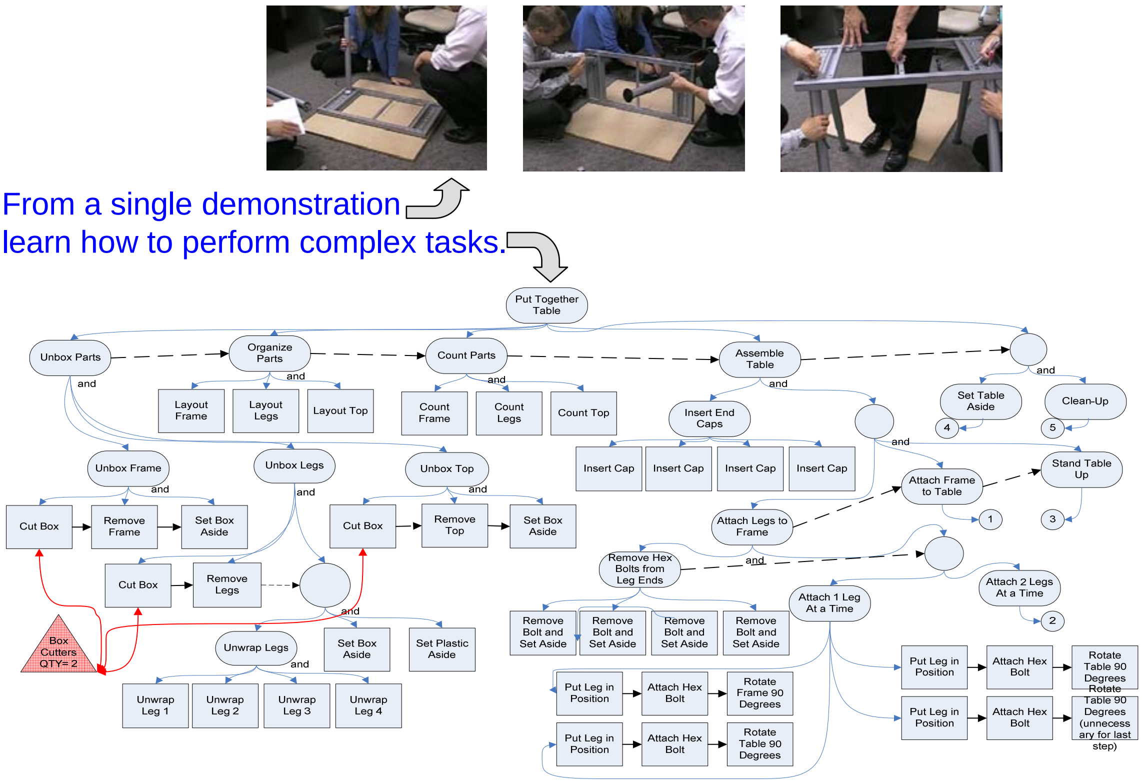

Figure 2 - Learn A Complex Plan From Single Set Of Observations

Once such a model is learned, it can be analyzed or incrementally extended by subsequent learning or by additional planning/reasoning activities. With such models systems can automate task performance or provide intelligent, contextually specific instruction to a human user. For instance, a person assembling an object might wind up in a dead-end state that does not lead to a goal. Using such a structure, an Integrated Learning system might back the human up to a prior state and then direct him/her down a different path to goal achievement.

For this program, no particular task model or plan construct is required, i.e., proposers must specify the construct they will learn. However, the model must have the richness that will enable a cognitive system to perform a meaningful or complex task with the model, i.e., flat sequences of steps are not of interest. Complex models have features like the following:

- Sequence, e.g., must do X, Y, Z, in an uninterrupted order.
- Precedence, e.g., must do X before Y.
- Temporal interactions, e.g., must do this X minutes before dong Y.
- Conditional performance, e.g., if (Z) then X else Y.
- Utility-driven choice, e.g., X is preferred over Y but both are ways to achieve task Z.
- Hierarchical organization – some elements are subtasks of others and items like conditional performance or sequencing can pertain to the higher level tasks.
- Joint action, e.g., two people must perform X together.

More details follow. Recall that each proposer must specify a complete Integrated Learner solution, including but not limited to (1) an Integrated Learning software framework, (2) an Integrated Learning Problem Solver, (3) reasoning components that will be included in the Integrated Learner, (4) simulation components that will be included in the Integrated Learner, (5) general world knowledge the learner will use, (6) domain knowledge the learner will use, (7) a cyberspace application, (8) a specification of the hierarchical plan or task models the learner will learn and produce, (9) the means of interacting with the human user, and (10) evaluation plans and metrics.

Concept Illustration

To illustrate the Integrated Learning concept, this section describes a *notional example* in which a simplified Integrated Learner learns physical assembly tasks. Specifically, the Integrated Learner produces a generalized plan of how to assemble a filing cabinet drawer by watching a person perform the assembly task just once.

<u>In contrast to this example, the Integrated Learning program will focus on learning generalized plans for processes completely contained in the computational world or cyberspace. For this program, the exact choice and articulation of an application is the responsibility of the proposer (see details and example applications below). We illustrate the Integrated Learning concept in a physical domain for illustrative purposes only.</u>

A simplified Integrated Learner is shown in Figure 3. This Integrated Learner incorporates two different types of domain knowledge (detailed geometric part models and topological constraints that describe how parts fit together), a simulation module of which the learner can ask "what if" questions about the physical world, and a domain planning reasoner from which the learner can ask for hypotheses about how things might be put together.

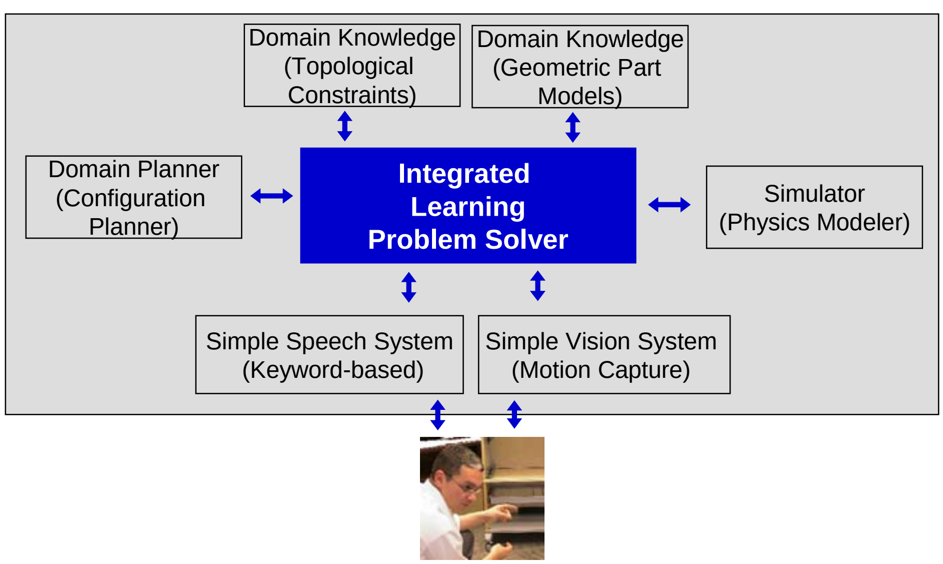

Figure 3 - A Simple Integrated Learner

Figure 4 shows a high-level control-flow view of how the simple Integrated Learner is used and operates. Generally, the person first performs the task for the learner, it learns a model, and then it uses the model to guide future human performance, e.g., providing contextually specific instruction by looking at the person's current state and the path he/she needs to take to reach a goal. Note that the system can also use the model to back the person out of a non-goal-reaching state to a different state from which the goal can be reached. For the Integrated Learning program, proposed solutions should address both the input side and the output side (how learned knowledge might be used) though the primary focus is on computer learning technologies (middle box in Figure 4).

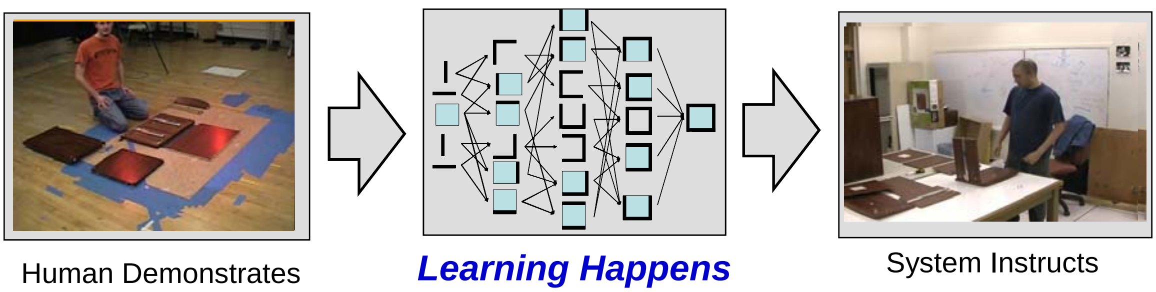

Figure 4 - High-Level Flow View of System Use

Figure 5 shows the first step in this process where the human demonstrates filing cabinet assembly for the system. In this example, the learner's input system records (1) spoken voice annotations, and (2) a "marker movie" of the tagged filing cabinet parts in motion.

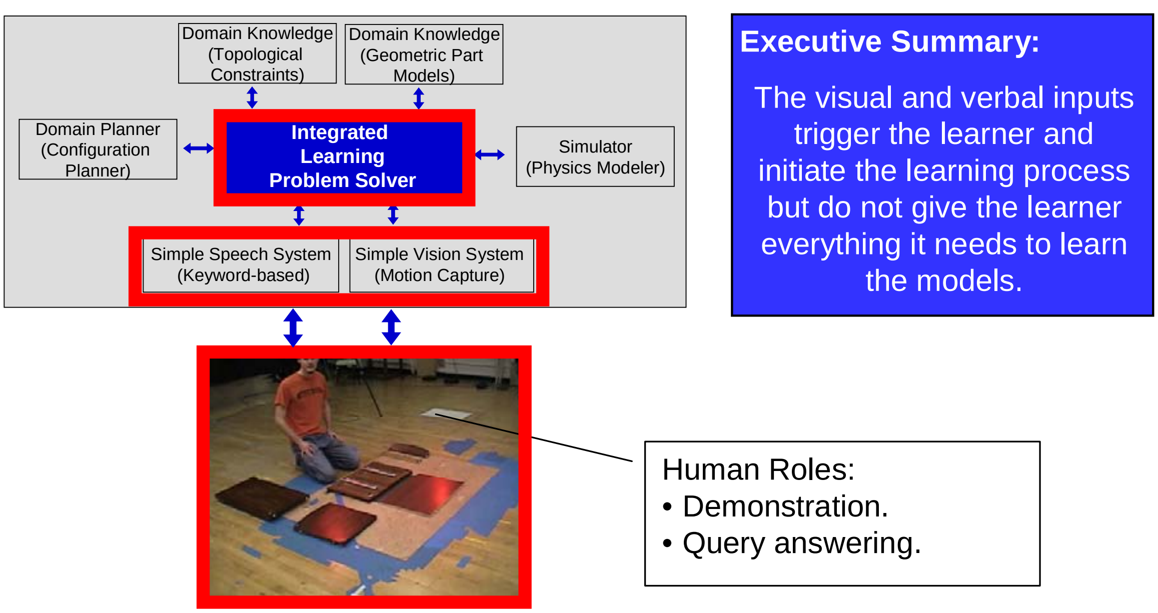

Figure 5 - Human Demonstrates for Learner

In response to this stimulus, the learner produces three learning goals: (1) figure out the exact sequence of steps the human performed, (2) generalize from this and attempt to generate more (different) ways to perform the assembly, and (3) if multiple different ways can be generated, then assess their difficulty.

In order to reason about plan steps the learner needs a more detailed model of the parts than is available from the marker movie. It might be possible for the learner to ask a human to hold each part in front of a camera so that it can learn the part models online. However, for this example we assume that the learner was given detailed part models *a priori* – conceptually the box of parts was opened and the parts described for the learner by a human before the assembly process began. In Figure 6, the learner uses the initial layout of the parts to match marked objects against parts that it knows about / has representations for in its domain knowledge. (Obviously, a reasoned matching scheme rather than an agreed upon initial layout would be a more general approach.)

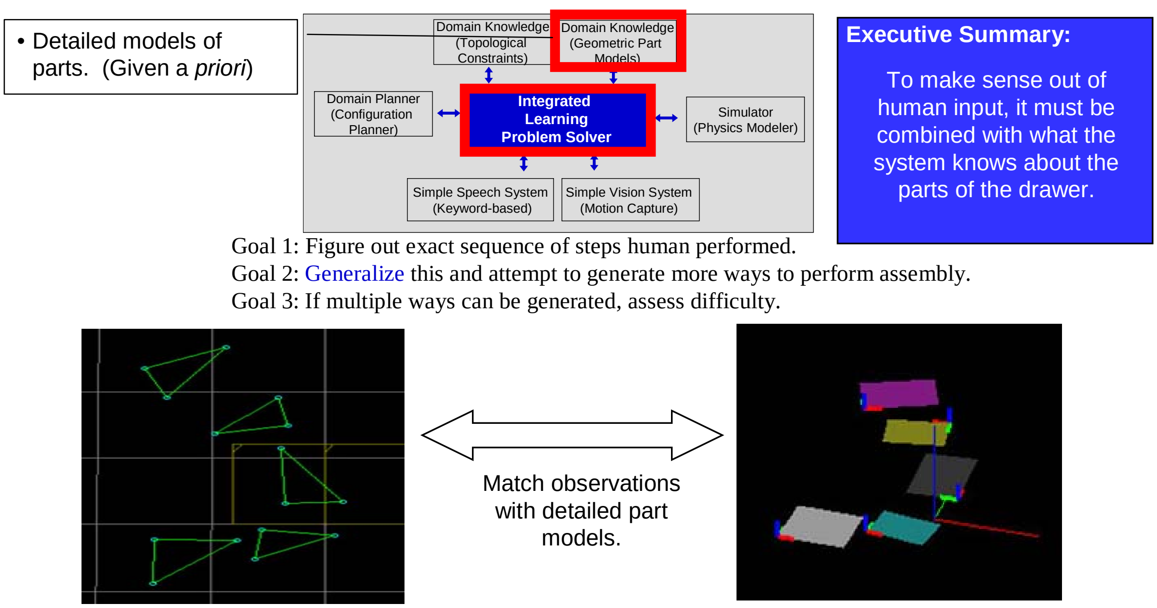

Figure 6 - Learner Correlates Domain Knowledge With Observations

The next step in the learning process, Figure 7, is to process the input and prepare it for detailed processing/learning. In order to do this, the learner combines the voice annotations (“step ended”) with the (noisy and obscured) marker movie to break the movie into segments that pertain to assembly steps. It is important to note that the learner could not simply generate an exact sequence of steps that were performed from the visual input – the input is far too uncertain, noisy, and incomplete. Instead the learner must combine these coarse observations with domain knowledge and domain reasoning in order to produce a generalized plan. Note also that finding the exact sequence of steps would not satisfy the goal of generalizing / finding alternative ways to perform the task and that the simple sequence itself does not contain the complexity desired for this program, i.e., creating a perfect input system is not the solution to this problem.

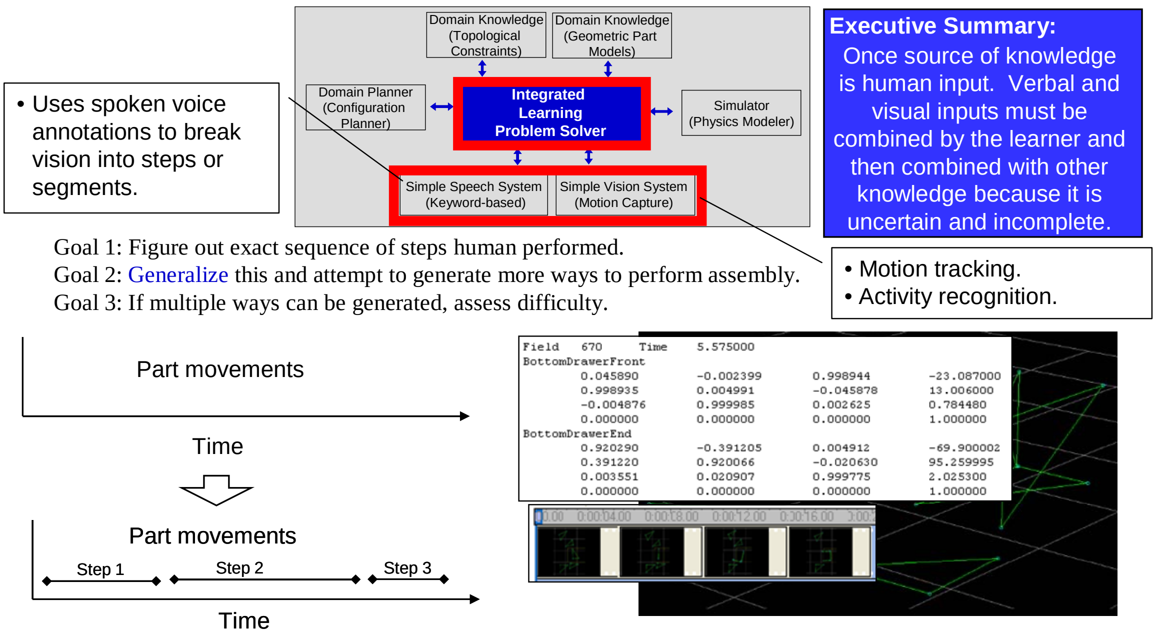

Figure 7 - Learner Processes Input for Learning

Now the learner must use its better internal models in conjunction with the input data / observations to try to make sense out of what it saw. This process is shown in Figure 8. Because the learner cannot make sense out of the marker data during any given step (because the parts are obscured), it instead looks at the final configuration of the parts at the end of each step or segment. In other words, rather than trying to extract exact movements from the noisy and obscured data, the learner looks for islands of stability (end states) and will then use other means (reasoning) to figure out how the final end states came about.

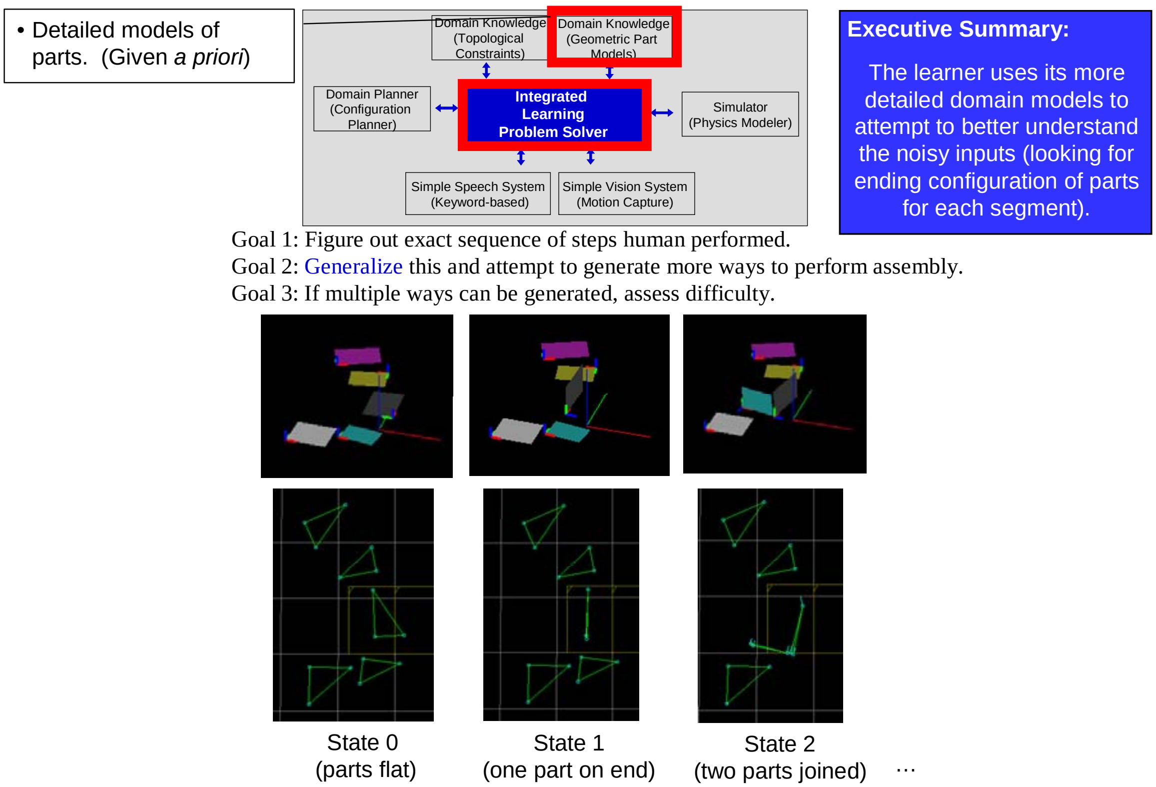

Figure 8 - Learner Tries To Determine Configuration End States From Uncertain Data

In the next step, Figure 9, the learner takes the end states that it tabulated and combines them with the geometric part models (detailed domain knowledge) and gives them to a domain planner that can look at the parts, look at the end states, and hypothesize a way to assemble the drawer. The network pictured on the bottom of the figure is the planner's output. The rightmost state is an end state that models the fully assembled drawer (also pictured) where it has four sides and a base. The internal states represent intermediate assembly points and the leftmost states are possible start states. Note that this structure may be incomplete (missing states) and may contain invalid states. This is because it was generated from end states that were based on noisy data and part models (generated by humans) that may be incomplete, contain errors, or not contain enough detail to detect certain issues. Essentially, at this point the learner has a hypothesized an uncertain way to put together the drawer.

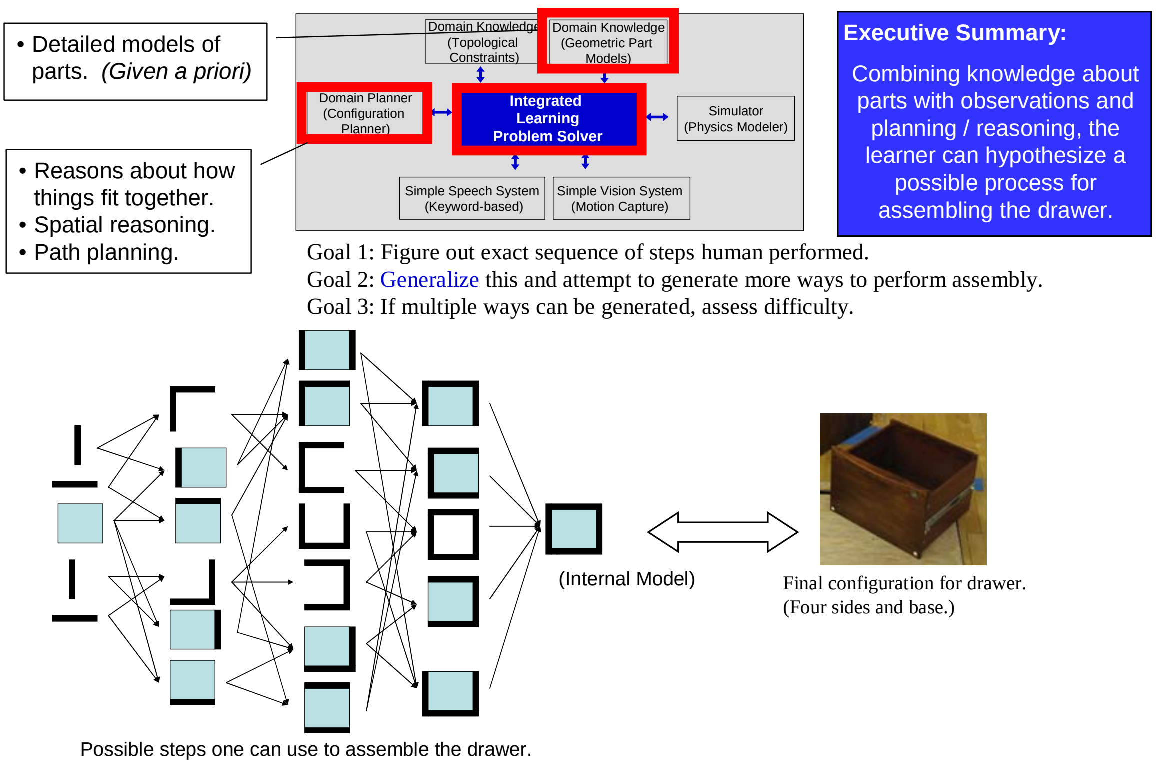

Figure 9 - Learner Produces Hypothesized Model Via Reasoning, Domain Knowledge, Observations, And Its Own Intermediate Results

The learner must then try to improve the hypothesized model. In Figure 10, the learner does this by walking the structure and trying to ascertain which transitions are valid, i.e., whether it is actually possible to get from one assembly state to another state given the way the parts fit together and the spatial constraints. To do this, for each transition, the learner passes the domain reasoner the transition, topological constraint knowledge (how parts fit together), and the detailed part models. It then asks if a plan that corresponds to the transition can be computed given the parts and spatial constraints involved. The figure shows the detail for two of penultimate states. Because a plan can be found for the uppermost transition, the learner knows it is a valid or legal transition. The planner is able to find the transition because it is possible to go from a state where the drawer has three sides and a base to a state where the drawer has four sides and a base, i.e., the face of the drawer is simply added. In contrast, the planner is unable to find a plan for the lower transition, i.e., it is not possible to put the drawer base on once the four sides have been joined because the base fits into a groove present in the four sides. Accordingly, the learner prunes the dead-end state from the output graph so that when the system guides a human through the process it can avoid states from which the human must backtrack. The learner should actually keep these states around but annotate the transitions so that if a human is already in a dead-end state when the system is brought online it can bring him/her to a prior state by reasoning about the process.

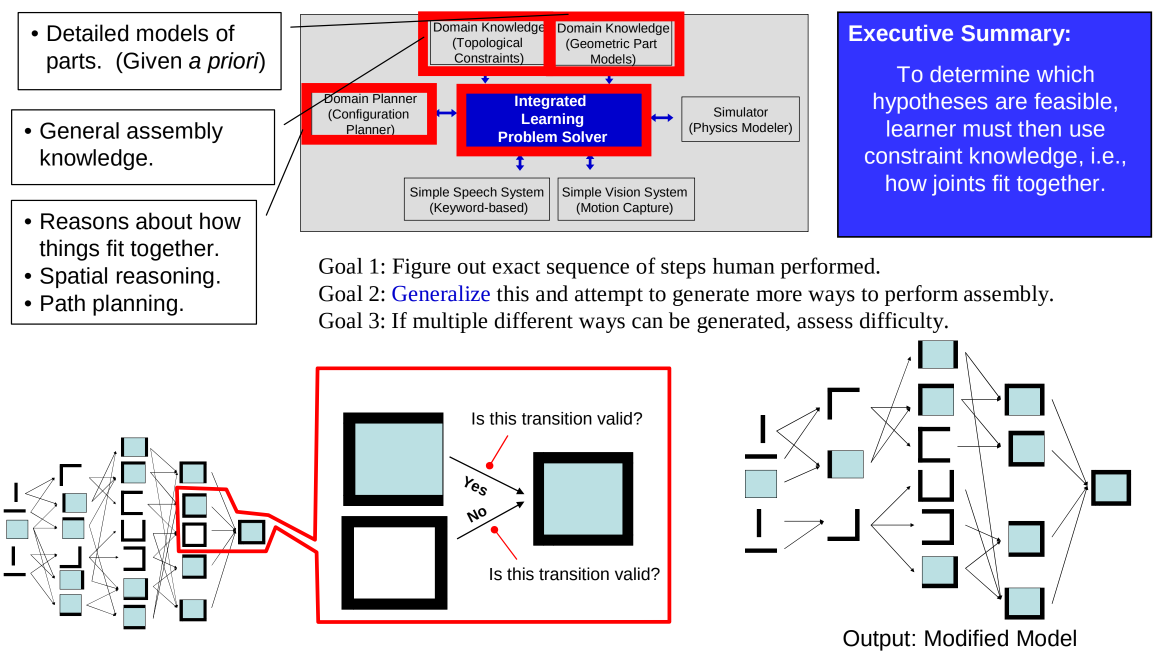

Figure 10 - Learner Refines The Model By Checking Each Transition Using Domain Knowledge About Parts And Constraints And Using Domain Reasoning

Recall that the learner has three learning goals, none of which have been achieved yet. At this point the learner has a hypothesized model for assembling the drawer. It must then go back over the model and verify that the end states it collected from the observations are actually contained in the model, i.e., double check to make sure what it thinks it saw is part of the model. This is shown in Figure 11. Because the learner is able to find the desired path (annotated with the red bounding box) we can consider the first goal achieved. Similarly, by combining reasoning with domain knowledge and the observations, the learner was able to generate a more general (multiple assembly approaches) way to assemble the drawer thus we can also consider the second goal achieved. Note that in the Integrated Learning program the learner should keep track of the uncertainties of these different states, e.g., the path with the red bounding box should be more certain than the other states in the model because corroborating observational input supports it though it may not be 100% certain at this point. In contrast, the other states and transitions were generated with uncertain data and at this point no corroborating information has been found. In a more sophisticated Integrated Learner, the learner may also be able to find corroborating information from other internal resources or from other learners – the process is not dependent entirely on observational inputs.

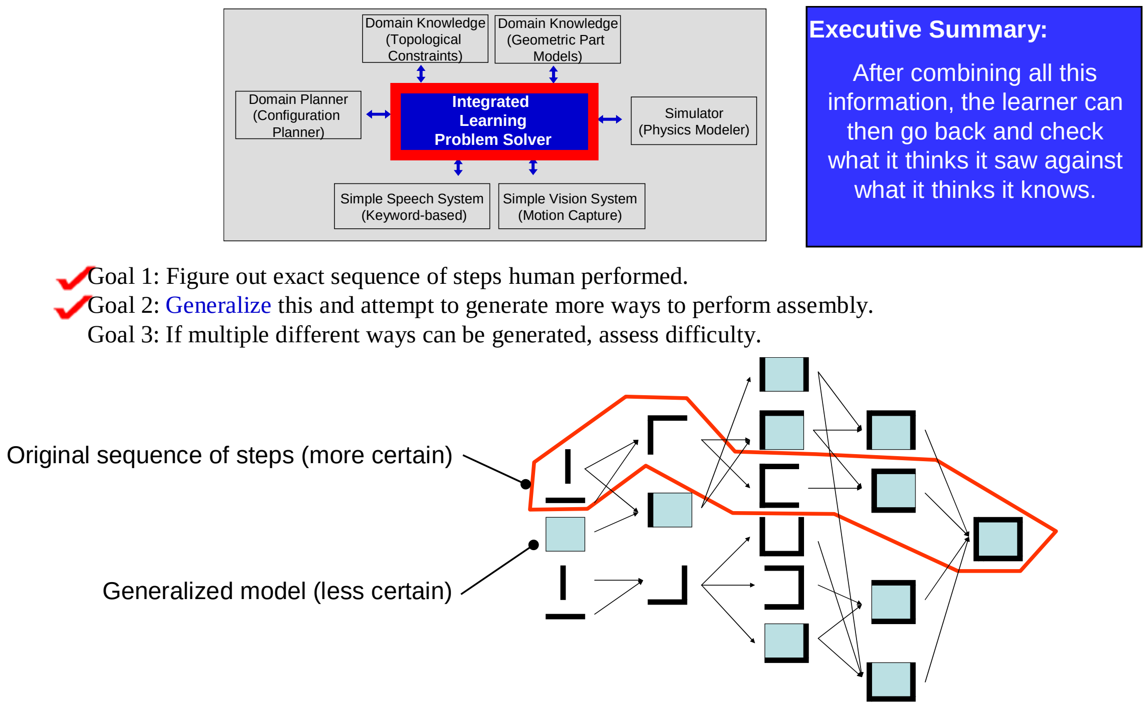

Figure 11 - Learner Has Produced A Plan That Matches Observations And A More Generalized Plan

The learner still has one goal remaining – assess the difficulty of the different possible paths through the network. In order to accomplish this, the learner uses a physics simulator to ask what-if questions about configuration stability, as shown in Figure 12. In this particular case, the metric being computed is the number of hands required to stabilize a part configuration while it is being moved / oriented for the next step in the process. To assess this, the learner constructs detailed part models that correspond to each state in the graph and then runs the part models through the physics simulator multiple times. In each run different aspects of the configuration are held stable while the rest are “released” (and gravity plays its role). The configuration shown requires three parts to be held stable during movement because the base and both sides must be held (the base does not screw in to the sides but rests in a groove). The transition out of the state is accordingly annotated with a "3" meaning that three hands are required. The learner walks the entire graph and performs a similar assessment for each transition. The final output is the annotated graph as shown. At this point, the learner has achieved all three goals and has a structure that is ready to use.

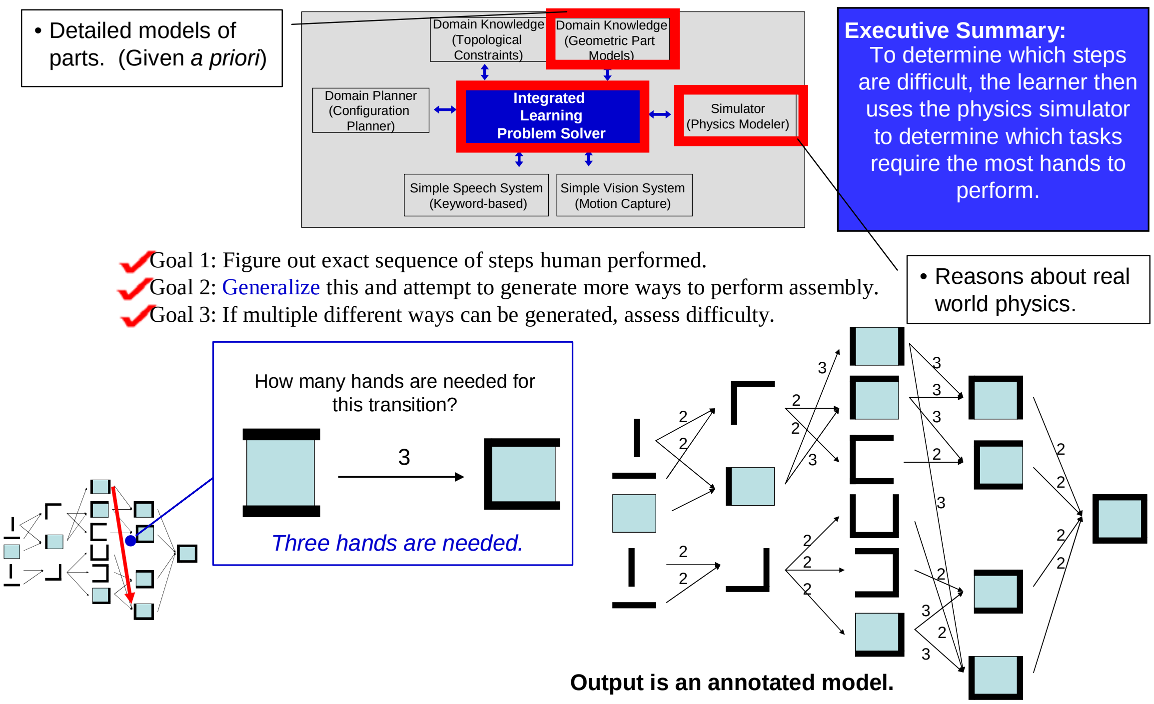

Figure 12 - Learner Uses Simulation To Improve Its Model / Figure Out Which Paths Are More Difficult

While the example ends here, in the program, Integrated Learners should view learning as a continuous, opportunistic, and incremental process. For instance, when new observational data comes in the learner could refine its model by adding or removing states and transitions. Uncertainties associated with different transitions or states could also be decreased/resolved (or increased if new data calls into question previous results). The learner may also generate new information by walking a person through the process, e.g., taking a person through a path and having them indicate they completed the assembly process should increase the certainty along that path. In some domains it may also be reasonable for the learner to guide a human through an uncertain assembly path in order to resolve or improve the uncertainties associated with the path, e.g., the person may provide feedback indicating that a transition the learner believes is possible is actually infeasible. In response to this feedback, the learner could generate a new learning goal to understand why it had the transition wrong and improve its own knowledge / reasoning / simulation components.

Note that this example does not demonstrate the learner computing a hierarchical model, i.e., the drawer is part of a filing cabinet. In fact, there are actually two drawers and the learner should recognize that they are both instances of the same task and they are subtasks of the filing cabinet assembly process. Note also that the structure generated is actually fairly simple relative to the general space of hierarchical plan representations. Similarly, the Learning Problem Solver in this example is very simple and control flow is just a sequential process. For the program, Integrated Learners must employ more sophisticated methods in control, knowledge assembly, representation, etc.

In this example, the Integrated Learner learned a general plan of how to perform a physical task not by receiving large numbers of training instances but by combining coarse observational data with domain knowledge (part models and topological constraints), reasoning (configuration planning), and simulations (asking what-if questions of a physics modeling system). The two most important aspects of this example are (1) how the Integrated Learner treats learning as an *integrated problem solving* process where learning is accomplished by combining observations (limited amounts and possibly uncertain) with knowledge and reasoning, and (2) that the Integrated Learner learns a generalized plan for assembling the object.

<u>In contrast to the physical assembly example, the Integrated Learning program will focus on learning generalized plans for processes in the computational world or cyberspace. For this program, the exact choice and articulation of an application is the responsibility of the proposers. Military domains are of interest but civilian application ideas are also welcome. More information and sample application domains appear in the following sections.</u>

Important Technical Ideas

Proposers are encouraged to identify their own important technical issues beyond the following list. However, solutions should:

- *Formulate learning as integrated problem solving.* Learners in this program must view learning as an integrated problem solving process where the learner reasons explicitly about what it knows, what it doesn't know, and what it needs to know in order to learn. The learner has explicit learning goals, formulates plans to achieve these goals, forms hypothesis where appropriate, tracks sources of uncertainty (possibly both symbolic and quantitative), and works to resolve these uncertainties. Control in an Integrated Learner should be both top-down / process driven and bottom-up / data-driven / opportunistic. Other advanced reasoning ideas are welcome. The key is to view learning as a problem solving process.

- *Flexibly combine different types of knowledge and reasoning.* Integrated Learners must be able to *assemble* information from many different sources including world knowledge, domain knowledge, reasoning, and simulation as needed. In a static environment where the class of plans being learned is fixed and the learner's information resources and their characteristics are known and fixed, it may be possible to "hard wire" the process of fitting together pieces of information in order to learn. In this program we are interested in learners that are flexible and dynamic – proposed solutions must be able to flexibly combine information from many different sources in order to learn by reasoning about how the information fits together.

- *Be open and extensible.* The use of “open” and “extensible” here means more than that the underlying software framework must be able to dynamically add components. The learning process itself must be able to dynamically take advantage of new sources of knowledge and new components for reasoning/simulation. Integrated Learners should not need a human to recode control algorithms, knowledge structures, etc., in order to incorporate new resources in their learning process. Conceptually, a user should be able to pass the learner new knowledge repositories, new reasoners, new simulation tools, etc., and have the learner incorporate the new information directly in its learning process.

- *Manage credit/blame assignment.* When different components within an Integrated Learner are able to provide information on the same topic, it is possible for inconsistencies to arise. For instance, one reasoning component might say “x cannot follow y because …” whereas another reasoning component might say “x must follow y because …” When such inconsistencies occur, the learner must either find evidence/information that proves or disproves one of those statements. However, if the learner is able to learn or be told (or some combination of these), which components are likely to be right or wrong in a given circumstance, it can better manage inconsistencies. More sophisticated ideas along these lines are welcome, e.g., credit/blame assignment could be combined with a deconfliction reasoning tool that could compare the “arguments” produced by different components.

- *Reason about the cost and value of information.* Integrated Learners will often have several options for obtaining needed information, e.g., the Learning Problem Solver might be able to query multiple reasoning systems to get information and might also be able to ask another Integrated Learner what it knows and might also be able to query a human. Each of these options may have cost/value trade-offs and the learner should reason explicitly about these when pursuing a course of action. For instance, if a quick-and-dirty heuristic will produce the right result 75% of the time the learner should try that before asking the human a question. Note that it may be appropriate to enrich this line of thought so that the learner treats “costs” and “values” as sets of attributes, e.g., “value” might entail (1) importance of information to achieving some learning goal, (2) importance of that goal, (3) timeliness, etc.

Sample Cyberspace Application Domains

<u>In contrast to the physical assembly example used here, the Integrated Learning program will focus on learning generalized plans for processes in the computational world or cyberspace. For this program, the exact choice and articulation of an application is the responsibility of the proposers. Military domains are of interest but civilian application ideas are also welcome.</u>

As with the physical world example, the program focus is on problems for which there is an underlying plan or process model that can be identified and learned by combining observations with knowledge, reasoning, and simulation. Accordingly, the domain should be such that reasoning and simulation technologies either already exist or can be reasonably constructed. Selecting domains that have a *constructive* element may be a good starting point. Domains in which forming or executing processes or plans is not a major factor should not be considered. Note that there are multiple classes of tasks performed with a computer and these different classes may have different characteristics. Example classes include: (1) tasks that are purely computational like complex data analysis or online procurement procedures, (2) tasks that incorporate reasoning about the physical world such as truck delivery planning, military mission planning, and air-tasking-order (ATO) generation, and (3) tasks that are essentially cyberspace versions of physical systems, e.g., CAD/CAM, microchip design, or circuit layout. Any of these classes may be appropriate. The cyberspace emphasis enables DARPA to focus directly on the Integrated Learning research and more easily obtain inputs (observational data) and provide meaningful output (tutorial or decision support to humans).

Sketches of several example applications follow:

1. Air Tasking Order (ATO) Planning – Create a system that learns human planning processes for Air Tasking Order (ATO) generation. Air operations campaign planning, which is carried out in the Air Operations Center (AOC), is a complicated time-sensitive process that currently involves hundreds of people using many different software systems to plan the activities of hundreds or thousands of aircraft, crews, support staff, support logistics, etc. The AOC is organized into cells, each with a specialized role. For instance the Strategy Cell identifies opponent's centers of gravity, updates rules of engagement, and coordinates with ground and naval commanders. The Combat Plans cell defines specific missions for individual aircraft by considering aircraft/team availability, the list of prioritized targets (for air strikes), and the list of prioritized collection goals (for air reconnaissance). The Air Tasking Order (ATO) is the primary product of the Combat Plans cell. The ATO provides detailed information about scheduled missions, mission plans, participants, goals, schedules, resource assignments, and contingencies for a 24-hour duty cycle. When forming the ATO, the human planners interact with many different software tools that perform functions like scheduling analysis, target / resource analysis, reconnaissance goal / resource analysis, airspace deconfliction, and logistics analysis. These tools are part of a suite called the Theater Battle Management Core System (TBMCS). The ATO planning process is complex due to the number of tools, humans, and pieces of information that must be consulted or factored-in. In this application, the Integrated Learner must learn the human ATO planning process or a significant and well defined portion of the planning process. Observational inputs to the Integrated Learner include, but are not limited to, the tools the human planner uses, the TBMCS screens he/she views, the targets identified by command, the mission objectives, and the priority order of targets and objectives. Other information an Integrated Learner might infer or gather includes the data state before and after a tool is invoked, the motivation for why the person used a particular tool at a particular time or in response to a particular situation (possibly obtaining this via voice annotation), and information about which goal a particular set of actions pertains to, e.g., "these three steps related to achieving mission goal X." Reasoning and simulation components the Integrated Learner may use include scheduling systems, airspace deconfliction, effects based analysis tools,

and adversarial simulation. World knowledge the Integrated Learner may use includes region specific geographic features, e.g., mountainous terrain, and weather. Domain knowledge the Integrated Learner might use includes aircraft armament options, aircraft fuel capacities, and maximum flight speed. By learning the ATO planning process the system should be able to (1) provide support to human planners engaged in similar missions, and (2) generalize to provide planning support to planners engaged in related missions.

2. Delivery AtlasQuest – Create a system that learns to construct plans for truck deliveries in urban areas. Factors a human planner might consider include simple distance and routes but also items like traffic patterns, tunnel height, streets that go one way during a particular time of day, regular one ways, traffic circles that are difficult for large trucks to traverse, availability of loading/unloading zones, congestion patterns (may be dependent on time of day), delivery route optimization, and even loading/unloading ordering (e.g., the glass goods must be loaded and dropped off either before or after the pianos but not interleaved), etc. The observational inputs to the Integrated Learner in this application are the route that a human planner has chosen, the time of day for which the route was planned, and possibly the loading/unloading of the truck. Information the learner might request, gather, or attempt to infer includes the current weather conditions, the truck identifier, the truck's dimensions, and the truck's cargo. Using this information the learner must infer the factors that the person considered, replicate the plan generated by the person, and then generalize from this plan so the system can proactively (1) find alternate plans to the same destinations for different times of day or that enable the driver to take a different route if traffic isn't flowing or detours occur, (2) find plans like this for other geographic destinations within the same city, and (3) find plans for the same or different destinations for different truck configurations. Reasoning tools an Integrated Learner might use idea include physical constraint reasoning (e.g., truck dimensions versus tunnel dimensions, turning radius), path planning, and cargo loading/capacities. Simulations the learner might run include time-of-day traffic modeling and weather modeling. Domain knowledge the learner might use includes regional maps and facts about truck delivery, e.g., for a truck to unload it must either back up to the building or its side doors must be facing the building. World knowledge the learner might use includes facts about how weather impacts traffic, e.g., rain causes congestion and snow causes even greater congestion. After learning these plans the system could then provide direct support to human planners and suggest routes.

3. CAD for Mechanical Assembly Planning – In this application, the Integrated Learner would learn how to create mechanical assembly plans, i.e., plans for how to put together a physical object that has been designed by another process. Factors a human planner might consider include minimizing the total time required for the assembly, minimizing the degree of difficulty of each step in the assembly, stability of the partially-assembled subparts, minimizing the risk of damage to the parts being assembled or to the surrounding environment, and constraints on the number of people and tools available during the assembly process. The human planner might also consider which paths are easier to encode using the design tool. The observational inputs to the system are the CAD models of the parts, the sequence of user interface commands that the person uses

to enter the assembly plan, and possibly natural language or spoken annotations such as "finished with the first drawer." Information the learner might request or infer includes the weight or material properties of parts and the resources (tools, machines, people) that will be available to use during the assembly process. From this, the learner must infer the factors the human planner considered, replicate the plan based on those factors, and then (1) generalize the plan so that the system can guide a new user through the process of building assembly plans for the same or similar objects (while allowing the new user to depart from the existing plan if feasible) and/or (2) use the plans to improve its own knowledge or reasoning components (e.g., formed plans could refine the learner's estimations of the stability of a given configuration, how likely a given part is to break, etc.). Reasoning tools the Integrated Learner might use include a configuration-space planner (check whether there exists a collision-free path to put a part into its desired position) and a grasping planner (check whether a person could grasp the object and manipulate it into position). Simulation tools the learner might use include a control synthesizer (synthesize the series of control forces that must be applied to the part to move it into the desired position), finite-element simulator (estimate risk of parts breaking), and a physical simulator (check for part stability, estimate number of people need to stabilize a part, lift a part, and put it into position). World knowledge the learner might use includes facts about how much weight a given person can easily lift and manipulate. Domain knowledge may include information about what materials each part is made of and the specific properties of the materials.

4. Bioinformatics – Create a system that learns to construct plans for gathering and processing information to support biologists that are exploring particular biological questions. Factors a human planner might consider include what information is available from what locations, how often the information is updated, how much it is trusted, how similar it is to the organism under study, what information is needed in order to access or filter the desired information, the quality of the information, and the amount of data. The observational inputs to the Integrated Learner in this application are the services that the human planner has chosen and the exact information flow between these services. Information the learner might request, gather, or attempt to infer includes the organism similarity, database update rates or freshness, service trust/level of curation, information quality, and alternate information and analysis sources. Using this information the learner must infer the factors that the person considered, replicate the plan generated by the person, and then generalize from this plan so the system can proactively (1) find alternate plans for processing different initial data, (2) find alternate plans upon failure of any component of the original plan at some future time, (3) suggest "better" plans, given the overwhelming number of resources available to biologists that increase almost daily. Reasoning tools an Integrated Learner might use include inferring organism similarity from phylogenic/taxonometric data, inferring information quality from linkage data or verification via multiple sources, and user preference modeling. Domain knowledge the learner might use includes the semantic data types used in genomics/proteomics: genomic sequences, expressed sequences, protein sequences, domains, motifs, etc. World knowledge the learner might use includes facts about service availability and network traffic.

5. Electro-Mechanical Devices Design (CAD) - The design of high-performance electro-mechanical products is a multi-step and iterative process that strives to balance tradeoffs between a large number of competing requirements, such as minimizing product cost, weight, and volume, while maximizing performance, ease-of-assembly, reliability, etc. To create a design, engineers use a suite of CAD tools and physical simulators that enable them to more quickly evaluate the implications and trade-offs of design choices along many dimensions, e.g., weight, volume, performance, and the manufacturing cost. Experienced engineers often have approximate design structures or plans that they follow in order to quickly hone in upon a relatively small set of key trade-offs, that are then explored in some detail. In contrast, novices often spend a great deal of time exploring the space of possible designs because they lack the internal plan models that enable them to rapidly identify key decision variables. In this application domain, the Integrated Learner would observe the design choices made by experts as they interact with a CAD system and infer the structure of their design process. Observational inputs would include layout decisions, choices of components and materials, and the sequence in which other software applications or simulation environments are used. The learner might request or gather information about the goals and requirements of the design and then attempt to infer motivation for a given action, e.g., a component was chosen to improve performance, a physical simulator was run to evaluate performance, or a layout was changed to ease manufacture. Reasoning tools an Integrated Learner might use include physical constraint reasoning. Existing simulation tools for modeling performance and assembly might be all that are required to inform the learners in this effort. Additional world and domain knowledge would include the relative cost and availability of parts, their performance characteristics, assembly options, and important design considerations, e.g., larger designs offer more convenient access to components, which eases maintenance. After learning the plans and strategies for using them, the system could guide novice engineers, make suggestions for initial designs, and facilitate rapid focus on key decision variables.

6. eScience Workflow Management – Create a system that learns plans for scientific workflow management. In disciplines like geology, scientists often have large data sets that are in remote repositories. They engage in long, many-stepped workflow plans in order to process the data. There are generally multiple alternative tools that can be used to process the data with each having different performance characteristics. The details of this general thrust are dependent on the specific domain, i.e., geologists will have different processing tools than oceanographers, but the general concept of watching an expert manage the workflow and learning the workflow, including conditional branching, applies across the domains.

7. Intelligent Travel Assistant (ITA) – Create a system that learns complex travel plans by watching a person plan his/her own travel. In contrast to conventional online systems in use today, the ITA must learn plans that involve a full spectrum of activities, e.g., flight arrangements, car rental, entertainment, dining reservations, sporting events, etc. Factors that a person might consider when planning include climate, weather, flight time preferences, flight length preferences (e.g., non-stop versus two segment), costs, entertainment preferences (e.g., likes jazz clubs), dining preferences (style of restaurant

and style of food served), and so forth. The observational inputs to the learner in this application are the actions the person takes when interacting with online systems for travel, reservations, etc. From this, the learner must infer the factors that the person considered when making arrangements, replicate the plan generated by the person, and then generalize the plan so the system can proactively (1) create alternate travel plans for the same dates but to a different destination, (2) generate alternate travel plans for the same destination/date range but select different options, e.g., pick different restaurants and book alternate entertainment events. Reasoning tools an Integrated Learner might use include path and travel-time planning. Simulations the learner might run include flight delay modeling to predict the implications of a missed flight. Domain knowledge the learner might use includes a mapping from venue to sport type (e.g., Wrigley Field is for baseball). World knowledge the learner might use includes general information about time-of-year weather patterns across the world. After learning these plans the system could then provide direct support to human travel planners and suggest alternate trips or alternate versions of a trip already planned.

When defining the application domain, proposers must clearly identify/describe the following:

- The **observations** the Integrated Learner will obtain from the human user.   In the physical domain the observations were part movements and voice annotations.
- The **world knowledge** the Integrated Learner will incorporate in its learning. In the physical domain the world knowledge might be facts about gravity or the properties of glues/adhesives.
- The **domain knowledge** the Integrated Leaner will incorporate in its learning. In the physical application the domain knowledge includes detailed models of parts and topological constraints that describe how the parts might be put together.
- The **reasoning systems** the Integrated Learner will utilize during learning. In the physical application, the learner employed a configuration planner to (a) generate a hypothesized graph describing the assembly process, and (b) to determine which transitions within the graph were possible.
- The **simulation systems** the Integrated Learner will utilize. In the physical application the learner employed a physics modeler to enable it to reason about configuration stability.
- The **process** or plan **being learned** and a generalization of it that the learner will compute. This should include a **figure of a hypothetical plan** that might be produced by the system.
- How the generated **plan can be used in an application** (this is different from the evaluation process – see the Evaluation Section), e.g., in the Delivery AtlasQuest application, the resulting plan can be used to provide decision support to human planners.
- Other **appropriate information**.

<u>DARPA is interested in applications that both enable Integrated Learning research and are challenging for the learner. It is the responsibility of the proposer to select a cyberspace application domain, properly scope it, articulate it, and support it as a choice.</u>

Technical Tasks, Emphasis Areas, and Teaming

There is exactly one technical task in this program:

- **Task I – Integrated Learner.** In simple terms, selected contractors are to build an Integrated Learner and evaluate it under DARPA's direction. This means constructing the Integrated Learning framework (including shared knowledge structures), creating an Integrated Learning Problem Solver, and populating the Integrated Learner with domain knowledge, world knowledge, reasoning components, and simulation components. Proposals must also define and characterize an application domain, as specified above, and specify the task models or generalized plans that they intend to learn (and the features thereof). Proposers will also develop their own test or evaluation harness and conduct their own evaluations under the direction of DARPA and an evaluation team appointed by DARPA. The primary components of this process and an envisioned level of effort by phase appears in the Program Phases section (below).

The technical task may require expertise in areas including, but not limited to machine learning, planning, problem solving, reasoning, knowledge representation, flexible software frameworks, and multi-modal interfaces. *Teaming is encouraged.* Note that one possible model for building an Integrated Learner is to develop multiple technologies for use within one framework, e.g., experimenting with multiple types of control or reasoning for the Learning Problem Solver or creating a Learning Problem Solver that employs multiple different approaches.

Evaluation and Metrics

This program has three primary evaluation goals: (1) measure each Integrated Learner's ability to learn, (2) verify that models being constructed by the learner include the features specified for the particular program phase, and (3) measure the contribution of each component to learning.

Conventional statistical machine learning evaluations (accuracy versus training set size) are not particularly meaningful for this program. Instead learning will be measured by demonstrating a task for the learner and then systematically verifying that it learned what was shown. This "point testing" process will be used for both general learning assessment and model feature verification. The metrics that will be used for evaluation are

- *Coverage percentile* – how much was it able to correctly learn? For instance, if it is shown 40 steps in a plan, was it able to learn 39 correct steps? Sample instances of the coverage metric include (others might include testing for iteration, conditionals, etc.): (1) step coverage percentile: SCP = # steps learned correctly / # steps shown; (2) choice coverage percentile: CCP = # choices learned correctly / # choices shown.
- *Error rate* – how many errors occurred and what was learned? Here we tabulate errors. Sample instances include (1) step error percentile: SEP = # steps learned incorrectly / # steps shown; (2) choice error percentile: CEP = # choices learned incorrectly / # choices shown.
- *Goal achievement* – can the learner achieve the desired goal? If so, how well? Goal achievement cannot be directly determined by the above two metrics. A learner might correctly learn 40 out of 40 steps but then follow those 40 steps with 12 errors in which the

goal is undone. There are also degrees of goal achievement, e.g., in a physical domain one can easily imagine putting together an object and having a few parts left over or having the object be functional but wobbly. The same ideas apply in computational domains, e.g., data analysis can be less complete than desired. To measure how well a goal is achieved, we will categorize goal achievement levels and score 0-10.

To determine what a "good score" is, the Integrated Learners will be benchmarked against human performance, i.e., a human will be shown how to do a task once and then asked to perform the task. Task performance will be scored using the metrics listed above. The learners will then be scored on a basis of how close to human performance they are able to achieve. For Phase I, the performance goal (go/no-go) is to achieve 65% of human performance. (Specifics for all phases are in the following section.)

To conduct the evaluations, proposers will work with DARPA and the evaluation team to specify classes of plans that their Integrated Learner will learn. Proposers can then "practice" with members of those classes. For the "test" runs, the evaluation team will sample from the specified classes to create different / new instances and will assess Integrated Learner performance on those instances. The evaluation team may also select members from different but related classes of plans and assess Integrated Learner performance on said instances. The details are somewhat application dependent. In the physical domain, for instance, a specific class might be "table assembly," a specific practice instance might be a Wal-Mart four legged table, and a specific test instance might be an IKEA four legged table. Proposals should include details of appropriate versions of this general protocol for their Integrated Learner / domain / plans.

To evaluate the contribution of each component to learning, a *contribution percentage* metric will be computed using ablation studies. The details of this will be dependent on the internals of a given Integrated Learner. For any Integrated Learner, no individual underlying reasoning, simulation, or knowledge source should regularly contribute most of the solution. 40% or less contribution for a given component is a target contribution percentage to consider when framing the application / formulating an approach. (Integrated Learners must learn by combining knowledge and reasoning from different sources.)

Note that aspects of the evaluation approach may be interdependent with the proposed Integrated Learner, its application domain, and the structures being learned. For instance, in some domains there may exist a "gold standard" against which learned models can be compared and evaluated in an automated fashion (though the general version of this for complex structures is intractable, i.e., graph isomorphism). Proposers should clearly state their metrics and specify their evaluation plans regardless of whether they are the same as articulated above or different.

Program Phases and Schedule

The Integrated Learning program will have four 12-month phases. Only Phase I will be funded initially, however, proposers should address all four phases. The general trend across all phases is (1) increasing amounts of reasoning, simulation, and knowledge in the Integrated Learner, (2) increasing sophistication of the Integrated Learning Problem Solver, (3) increasing complexity desired of the learned models, and (4) increasing levels of proficiency are required. The phases and their focus areas are as follows:

- *Phase I (12 months) – Integrated Learning.*
  - *Description:* The core of an Integrated Learner will be constructed in this phase. The learner must integrate multiple different components for reasoning, simulation, and knowledge in order to learn. The learning problem solver may be basic but must still assemble knowledge from the different components in order to learn.
  - *Integration:* Integrated learner with three reasoners/simulators and two types of knowledge.
  - *Problem solving:* Simple learning problem solver.
  - *Task models:* Learn steps and hierarchies but not learning choices, resource relationships, or iteration.
  - *Go/No-go:* 65% of human performance.

- *Phase II (12 months) – Learning Problem Solving.*
  - *Description:* The Integrated Learning Problem Solver is enhanced in this phase to incorporate more sophisticated concepts of problem solving in order to learn. Note that the language below can apply to many different AI paradigms, e.g., planning + monitoring, blackboard systems, etc. Interesting and novel solutions to Learning Problem Solving are sought.
  - *Integration:* Add two more reasoners/simulators and two types of knowledge. Modify five of the reasoners/simulators/knowledge components so that they emit uncertainty information for the Learning Problem Solver (this information may be both quantitative and symbolic).
  - *Problem solving:* Modify the learning problem solver:
    - Has explicit goals.
    - Knows what it doesn't know and plans to find it out.
    - Tracks sources of uncertainty and decides how to resolve them.
    - Learns / modifies structures both opportunistically and incrementally.
    - Employs both process-driven (top down) and opportunistic (bottom up) control.
  - *Task models:* Must learn choices, must incrementally learn / add to structures.
  - *Go/No-go*: 85% of human performance.

- *Phase III (12 months) – Open The Learner, Cost/Benefits, Credit/Blame.*
  - *Description:* In this phase the Integrated Learner is "opened" so (a) it can incorporate new reasoning/simulation/knowledge components at run time and use them in the learning process, and (b) the roles of the underlying components can change as they also learn over time. The general objective is for the learner to dynamically and flexibly determine the roles of the different components in the learning process and dynamically/flexibly determine how to assemble knowledge from different components. (Some representation of what a component is, does, can do, etc., is probably one aspect of solving this technical problem.) For instance, referring back to the Integrated Learner that learned plans for filing cabinet assembly, said learner might initially be designed to utilize only a configuration planner and detailed domain knowledge. By "opening" the learner, a sophisticated user should be able to add the physics modeler and have the

system utilize this new capability without (1) reprogramming the learner's control flow / processing, or (2) reprogramming the learner's shared knowledge structures. In essence, the learner should register the new module / new capability and automatically take advantage of it when appropriate. In this phase the learner must also add technologies to manage credit/blame and perform cost/benefit analysis. Note that in this phase the learner will also be required to learn tasks involving multiple parties. (Some applications may not have a multi-person task analog. In these cases the proposer must clearly articulate another challenging problem for the learner to address.)
  - *Integration:* Add one new reasoner/simulation component and one new knowledge source. Add uncertainty representation to any remaining components for which this is appropriate. Open the system so it can dynamically, at run time, add any of the components and have them be properly incorporated in the problem solving process. (Performers may be asked to incorporate other components selected by DARPA and developed by third parties[1] – solutions must be open.). Adapt four of the reasoning/simulation/knowledge components so they change / learn online also. (In-the-small evaluations will be used to test this learning.)
  - *Problem solving:* Dynamically incorporates new reasoning/knowledge into the process, reasons about costs/benefits when deciding what to do next, learns credit/blame assignments.
  - *Task models:* Add the ability to learn multi-person task structures.
  - *Go/No-go:* 105% of human performance.[2]

- *Phase IV (12 months) – Learn Meta Processes, Expectation Driven Learning, Sharing*
  - *Description:* In this phase the learner must be modified to be able to "step away" from the details of the process it is learning and to learn general process or meta process knowledge. In the physical domain, this would entail learning that the first phase in assembling an object is generally to lay out and count the parts and the last phase is often to tighten the bolts. This learned information must be fed back into the Learning Problem Solver so the meta process information guides/improves learning, i.e., provides expectation-driven learning. In this phase the learner must also be opened even further so it can share information (low-level data, mid-level hypothesis, and high-level conclusions) with other learners.
  - *New:* Learn general assembly knowledge, e.g., final step is often to tighten all bolts.
  - *Integration:* Different Integrated Learners must share information in order to learn.
  - *Problem solving:* Use learned assembly knowledge as expectations in the learning process. Add the ability to incorporate new information from different learners.
  - *Task models:* Learn resource relationships and iteration.
  - *Go/No-go:* 125% of human performance.[1]

A preliminary program schedule is shown in Figure 13. Not all of the detailed milestones may be appropriate for every proposed solution. Whole artifact evaluations and ablation studies will

[1] In which case a common API for such components will be specified in the first quarter of Phase III.
[2] If a ceiling effect occurs during evaluation and the humans are able to score 100%, the problem difficulty will be increased. In the unlikely event that humans always score 100% on a given application, learning progress will be measured by constructing a new test and comparing the performance on the previous phase to the current phase.

be conducted at the end of each phase as indicated and the general level of effort information should be considered when proposers construct their own schedules and plans.

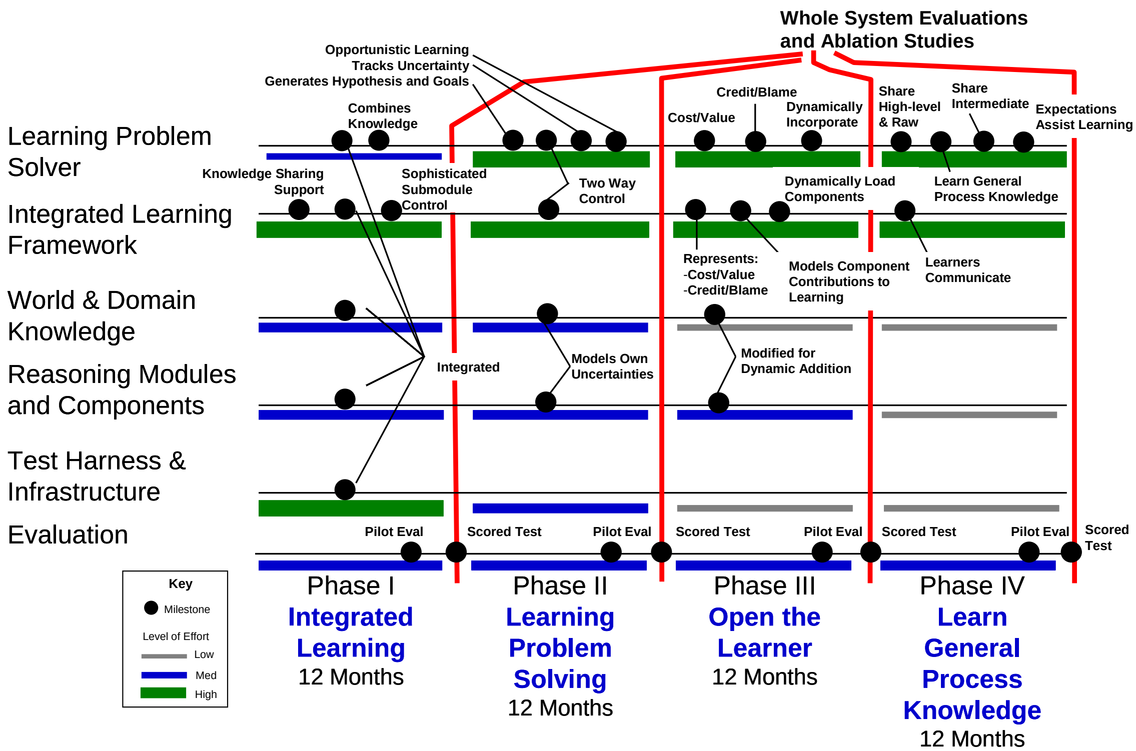

Figure 13 – Preliminary Program Schedule (Proposer's Milestones May Differ But General Trends And Evaluation Timing Apply)

PROGRAM SCOPE

Proposed research should investigate innovative approaches and techniques that lead to or enable revolutionary advances in the state-of-the-art. Proposals are not limited to the specific strategies listed above, and alternative visions will be considered. However, proposals should be for research that substantially contributes towards the goals stated. Specifically excluded is research that primarily results in minor evolutionary improvement to the existing state of practice or focuses on special-purpose systems or narrow applications.

GENERAL INFORMATION

Proposals not meeting the format described in this pamphlet may not be reviewed. Proposals **MUST NOT** be submitted by fax or e-mail; any so sent will be disregarded. This notice, in conjunction with the BAA 05-43 FBO Announcement and all references, constitutes the total BAA. A Frequently Asked Questions (FAQ) list may be provided. The URL for the FAQ will be specified on the DARPA/IPTO BAA Solicitation page. No additional information is available, nor will a formal Request for Proposal (RFP) or other solicitation regarding this announcement be issued. Requests for same will be disregarded. All responsible sources capable of satisfying the Government's needs may submit a proposal that shall be considered by DARPA. Historically Black Colleges and Universities (HBCUs), Small Disadvantaged Businesses and Minority Institutions (MIs) are encouraged to submit proposals and join others in submitting proposals. However, no portion of this BAA will be set aside for Small Disadvantaged Business, HBCU and MI participation due to the impracticality of reserving discrete or severable areas of this research for exclusive competition among these entities.

Proposals selected for funding are required to comply with provisions of the Common Rule (32 CFR 219) on the protection of human subjects in research (http://www.dtic.mil/biosys/downloads/32cfr219.pdf) and the Department of Defense Directive 3216.2 (http://www.dtic.mil/whs/directives/corres/html2/d32162x.htm). All proposals that involve the use of human subjects are required to include documentation of their ability to follow Federal guidelines for the protection of human subjects. This includes, but is not limited to, protocol approval mechanisms, approved Institutional Review Boards, and Federal Wide Assurances. These requirements are based on expected human use issues sometime during the entire length of the proposed effort.

For proposals involving "greater than minimal risk" to human subjects within the first year of the project, performers must provide evidence of protocol submission to a federally approved IRB *at the time of final proposal submission to DARPA.* For proposals that are forecasted to involve "greater than minimal risk" after the first year, a discussion on how and when the proposer will comply with submission to a federally approved IRB needs to be provided in the submission. More information on applicable federal regulations can be found at the Department of Health and Human Services – Office of Human Research Protections website (http://www.dhhs.gov/ohrp/).

DARPA has determined that work for this program is to be funded by budget category 6.2 (Applied Research). This means that research performed under this program on-campus at a university is considered contracted fundamental research; therefore, public releases of information about such research are not subject to prior Government review. The definition of CONTRACTED FUNDAMENTAL RESEARCH is contained in DOD Instruction 5230.27 and can be found at http://www.dtic.mil/whs/directives/corres/pdf2/i523027p.pdf. Public release of information about research performed under circumstances other than those described above is subject to prior government review, according to the procedures available at http://www.darpa.mil/tio.

Security classification guidance on a DD Form 254 (DoD Contract Security Classification Specification) will not be provided at this time since DARPA is soliciting ideas only. After reviewing incoming proposals, if a determination is made that contract award may result in access to classified information, a DD Form 254 will be issued upon contract award. **If you choose to submit a classified proposal you must first receive the permission of the Original Classification Authority to use their information in replying to this BAA.**

SUBMISSION PROCESS

This BAA requires completion of an online Cover Sheet for each Proposal prior to submission. To do so, the offeror must go to **http://www.dyncorp-is.com/BAA/index.asp?BAAid=05-43** and follow the instructions there. Each offeror is responsible for printing the BAA Confirmation Sheet and attaching it to every copy. The Confirmation Sheet should be the first page of the Proposal. If an offeror intends to submit more than one Proposal, a unique UserId and password must be used in creating each Cover Sheet. Failure to comply with these submission procedures may result in the submission not being evaluated.

Proposers must submit the original and **2** copies of the full proposal *and* **2** electronic copies (i.e., **2** separate disks) of the full proposal (in PDF or Microsoft Word 2000 for IBM-compatible format on a 3.5-inch floppy disk or cd). **Mac-formatted disks will not be accepted**. Each disk must be clearly labeled with BAA 05-43, proposer organization, proposal title (short title recommended) and "Copy ___ of **2**". The full proposal (original and designated number of hard and electronic copies) must be submitted in time to reach DARPA by 12:00 PM (ET) September 14, 2005, in order to be considered during the initial evaluation phase. However, BAA 05-43, Integrated Learning will remain open until 12:00 NOON (ET) July 11, 2006. Thus, proposals may be submitted at any time from issuance of this BAA through July 11, 2006. While the proposals submitted after the September 14, 2005, deadline will be evaluated by the Government, proposers should keep in mind that the likelihood of funding such proposals is less than for those proposals submitted in connection with the initial evaluation and award schedule. DARPA will acknowledge receipt of submissions and assign control numbers that should be used in all further correspondence regarding proposals.

Restrictive notices notwithstanding, proposals may be handled for administrative purposes by support contractors. These support contractors are prohibited from competition in DARPA technical research and are bound by appropriate non-disclosure requirements. Input on technical aspects of the proposals may be solicited by DARPA from non-Government consultants /experts who are also bound by appropriate non-disclosure requirements. However, non-Government technical consultants/experts will not have access to proposals that are labeled by their offerors as "Government Only". Use of non-government personnel is covered in FAR 37.203(d).

REPORTING REQUIREMENTS/PROCEDURES:

The Award Document for each proposal selected and funded will contain a mandatory requirement for submission of DARPA/IPTO Quarterly Status Reports and an Annual Project Summary Report. These reports, described below, will be electronically submitted by each awardee under this BAA via the DARPA/IPTO Technical – Financial Information Management System (T-FIMS). The T-FIMS URL will be furnished by the government upon award. Detailed data requirements can be found in the Data Item Description (DID) DI-MISC-81612A available on the Government's ASSIST database (http://assist.daps.dla.mil/quicksearch/).

PROPOSAL FORMAT

Proposals shall include the following sections, each starting on a new page (where a "page" is 8-1/2 by 11 inches with type not smaller than 12 point) and with text on one side only. The submission of other supporting materials along with the proposal is strongly discouraged. Sections I and II (excluding the submission cover/confirmation sheet and section N) of the proposal shall not exceed the total of the maximum page lengths for each section as shown in braces { } below.

Section I. Administrative

The BAA Confirmation Sheet {1 page} described under "Submission Process" will include the following:

A. BAA number;
B. Technical topic area;
C. Proposal title;
D. Technical point of contact including: name, telephone number, electronic mail address, fax (if available) and mailing address;
E. Administrative point of contact including: name, telephone number, electronic mail address, fax (if available) and mailing address;
F. Summary of the costs of the proposed research, including total base cost, estimates of base cost in each year of the effort, estimates of itemized options in each year of the effort, and cost sharing if relevant;
G. Contractor's type of business, selected from among the following categories: "WOMEN-OWNED LARGE BUSINESS," "OTHER LARGE BUSINESS," "SMALL DISADVANTAGED BUSINESS [*Identify ethnic group from among the following: Asian-Indian American, Asian-Pacific American, Black American, Hispanic American, Native American, or Other*]," "WOMEN-OWNED SMALL BUSINESS," "OTHER SMALL BUSINESS," "HBCU," "MI," "OTHER EDUCATIONAL," "OTHER NONPROFIT", or "FOREIGN CONCERN/ENTITY."

Section II. Detailed Proposal Information

This section provides the detailed discussion of the proposed work necessary to enable an in-depth review of the specific technical and managerial issues.

Page-counts are maximums.

A. {1 Page} Innovative claims for the proposed research.
This page is the centerpiece of the proposal and should succinctly describe the unique proposed contribution.

B. {1 Page} Proposal Summary
The summary provides a top-level view of the proposal. It contains a synopsis (or "sound byte") for each of the areas defined below. It is important to make the synopses as explicit and informative as possible. Where appropriate, the summary should also cross-reference the proposal page number(s) where each area is elaborated. The summary areas are:

1. Main goals of the proposed research (stated in terms of new, operational capabilities).

2. Tangible benefits to end users (i.e., benefits of the capabilities above).

3. Critical technical barriers or technical limitations that have, in the past, prevented the operational capabilities/benefits described above.

4. Summary of the cyberspace application being proposed and the hierarchical task or plan models that will be learned.

5. Main elements of the proposed Integrated Learner approach.

6. Summary of why the proposed approach will overcome the technical barriers.

7. Expected results of this work (unique/innovative/critical capabilities to result from this effort, and form in which they will be defined).

8. Evaluation plan summary.

9. Cost of the proposed effort for each performance year.

C. {2 Pages} Research Objectives:

1. Problem Description. Provide a concise description of the problem areas addressed by this Integrated Learning research. Make this specific to your Integrated Learning approach, application domain, and hierarchical task or plan models.

2. Research Goals. Identify specific research goals for your Integrated Learner approach. Goals can be both system level and detailed, e.g., pertaining to internal components.

3. Expected Impact. Describe expected impact of your Integrated Learning research to both the specific problem being solved and to the larger computer learning technology base.

D. Technical Approach and Evaluation:

1. {2 Pages} Application and Models. Provide a detailed description of the computational or cyberspace application domain being proposed for this work and a description of the hierarchical task or plan models the learner will produce.

2. {12 Pages} Technical Approach. Provide a detailed description of the technical approach being taken to create an Integrated Learner. This should include, but is not limited to, description of: (1) the Integrated Learning software framework, (2) the Integrated Learning Problem Solver, (3) reasoning components that will be included in or developed for the learner, (4) simulation components that will be included in or developed for the learner, (5) world knowledge the learner will use, (6) domain knowledge the learner will

use, and (7) the means of interacting with the human user or a description of where/how the learner will obtain its (limited) observational data.

3. {2 Pages} Comparison with Current Technology. Describe state-of-the-art approaches and the limitations that relate to the proposed Integrated Learning research.

4. {2 Pages} Evaluation/Experimentation Plans and Metrics. Proposers should clearly define appropriate metrics and evaluation plans for their approach. These plans may build directly on those given in the PIP though some customization may be appropriate. The program-wide evaluation goals are given in the Evaluation Section; proposers should address these goals at a minimum. Though the planned model is for each team to conduct its own evaluations under the direction of DARPA and an independent evaluation team, appointed by DARPA, proposers should also be willing to work with other contractors in order to develop joint experiments in a common testbed environment. Related to the program-wide evaluations, proposers should expect to participate in workshops to provide specific technical background information to DARPA, attend semi-annual Principal Investigator (PI) meetings, and participate in numerous other coordination meetings via teleconference or Video Teleconference (VTC). Funding to support these various group experimentation efforts should be included in the bid.

E. {3 Pages} Statement of Work (SOW). Provide a statement of work, written in plain English, outlining the scope of the effort and citing specific tasks to be performed, references to specific subcontractors if applicable, and specific contractor requirements.

F. Schedule and Milestones. This section should include:

1. {1 Page} Schedule Graphic. Provide a graphic representation of project schedule including detail down to the individual effort level. This should include but not be limited to, a multi-phase development plan, which demonstrates a clear understanding of the proposed research; and a plan for periodic and increasingly robust experiments over the project life that will show applicability to the overall program concept. Show all project milestones. Use absolute designations for all dates.

2. {3 Pages} Detailed Individual Effort Descriptions. Provide detailed task descriptions for each individual effort and/or subcontractor in schedule graphic.

G. {1 Page} Project Management and Interaction Plan. Describe the project management and interaction plans for the proposed work. If proposal includes subcontractors that are geographically distributed, clearly specify working / meeting models. Items to include in this category include software/code repositories, physical and virtual meeting plans, and online communication systems that may be used.

H. {2 Pages} Deliverables Description. List and provide detailed description for each proposed deliverable. Include in this section all proprietary claims to results, prototypes, or systems supporting and/or necessary for the use of the research, results, and/or prototype. If there are no proprietary claims, this should be stated. The offeror must submit a separate list of all technical data or computer software that will be furnished to the Government with other than unlimited rights (see DFARS 227.) Specify receiving organization and expected delivery date for each deliverable.

I. {1 Page} Technology Transition and Technology Transfer Targets and Plans. Discuss plans for technology transition and transfer. Identify specific military and commercial organizations for technology transition or transfer. Specify anticipated dates for transition or transfer.

J. {4 Pages} Personnel and Qualifications. List of key personnel, concise summary of their qualifications, and discussion of proposer's previous accomplishments and work in this or closely related research areas. Indicate the level of effort to be expended by each person during each contract year and other (current and proposed) major sources of support for them and/or commitments of their efforts. *DARPA expects all key personnel associated with a proposal to make substantial time commitment to the proposed activity.*

K. {1 Page} Facilities. Description of the facilities that would be used for the proposed effort. If any portion of the research is predicated upon the use of Government Owned Resources of any type, the offeror shall specifically identify the property or other resource required, the date the property or resource is required, the duration of the requirement, the source from which the resource is required, if known, and the impact on the research if the resource cannot be provided. If no Government Furnished Property is required for conduct of the proposed research, the proposal shall so state.

L. {2 Pages} Cost Summaries. This section shall contain two tables: (1) The first table must summarize the proposed costs but break them down by task and phase, i.e., show the costs of each task for each phase with the task labels on the y-axis and the four phases on the x-axis. It may be appropriate to create a subtotal under some closely related tasks. Table entries should contain the dollar figure and a percentage that specifies the percentage of that phase's total costs that are allocated to said task. (2) The second table should show the costs broken down by prime/subcontractor by phase, i.e., the labels of the prime/subcontractors should be on the y-axis and the four phases on the x-axis. Table entries should contain the dollar figure and a percentage that specifies the percentage of that phase's total costs allocated to said prime or subcontractor.

M. {5 pages} Cost Details. Cost proposals shall provide a detailed cost breakdown of all direct costs, including cost by task, with breakdown into accounting categories (labor, material, travel, computer, subcontracting costs, labor and overhead rates, and equipment), for the entire contract and for each **calendar year**, **divided into quarters**. Where the effort consists of multiple portions that could reasonably be partitioned for purposes of funding, these should be identified as contract options with separate cost estimates for each.

N. GFP. Contractors requiring the purchase of information technology (IT) resources as Government Furnished Property (GFP) **MUST** attach to the submitted proposals the following information:

1. A letter on Corporate letterhead signed by a senior corporate official and addressed to **Dr. Tom Wagner,** DARPA/IPTO, stating that you either can not or will not provide the information technology (IT) resources necessary to conduct the said research.

2. An explanation of the method of competitive acquisition or a sole source justification, as appropriate, for each IT resource item.

3. If the resource is leased, a lease purchase analysis clearly showing the reason for the lease decision.

4. The cost for each IT resource item.

O. Organizational Conflict of Interest. Awards made under this BAA may be subject to the provisions of the Federal Acquisition Regulation (FAR) Subpart 9.5, Organizational Conflict of Interest. All offerors and proposed subcontractors must affirmatively state whether they are supporting any DARPA technical office(s) through an active contract or subcontract. All affirmations must state which office(s) the offeror supports, and identify the prime contract number. Affirmations should be furnished at the time of proposal submission. All facts relevant to the existence or potential existence of organizational conflicts of interest, as that term is defined in FAR 2.101, must be disclosed, organized by task and year. This disclosure shall include a description of the action the Contractor has taken, or proposes to take, to avoid, neutralize, or mitigate such conflict.

IMPORTANT NOTE: IF THE OFFEROR DOES NOT COMPLY WITH THE ABOVE STATED REQUIREMENTS, THE PROPOSAL WILL BE REJECTED.

Section III. Additional Information

A bibliography of relevant technical papers and research notes (published and unpublished) that document the technical ideas, upon which the proposal is based, may be included in the proposal submission. Provide one set for the original full proposal and one set for each of the **x** full proposal hard copies. Please note: The materials provided in this section, and submitted with the proposal, will be considered for the reviewer's convenience only and not considered as part of the proposal for evaluation purposes.

EVALUATION AND FUNDING PROCESSES

Proposals will not be evaluated against each other, since they are not submitted in accordance with a common work statement. DARPA's intent is to review proposals as soon as possible after they arrive; however, proposals may be reviewed periodically for administrative reasons. For evaluation purposes, a proposal is the document described in PROPOSAL FORMAT Section I and Section II (see below). Other supporting or background materials submitted with the proposal will be considered for the reviewer's convenience only and not considered as part of the proposal.

Evaluation of proposals will be accomplished through a scientific review of each proposal using the following criteria, which are listed in descending order of relative importance:

(1) Overall Scientific and Technical Merit: The overall scientific and technical merit must be clearly identifiable and compelling. The technical concepts should be clearly defined and developed. The technical approach must be sufficiently detailed to support the proposed concepts and technical claims. Proposal must clearly define metrics and evaluation plans. Proposal must also clearly define system integration approach and plans.

(2) Innovative Technical Solution to the Problem: Offerors should apply new and/or existing technology in an innovative way that supports the objectives of the proposed effort. The proposed concepts and systems should show breadth of innovation across all the dimensions of the proposed solution.

(3) Offeror's Capabilities and Related Experience: The qualifications, capabilities, and demonstrated achievements of the proposed principals and other key personnel for the primary and subcontractor organizations must be clearly shown.

(4) Plans and Capability to Accomplish Technology Transition: The offeror should provide a clear strategy and plan for transition to military forces (and commercial sector, where applicable). Offerors should consider involving potential military transition partners, as appropriate, in any proposed experiments, tests and demonstrations. Offerors should also provide a plan for transition of appropriate technology components and information to the user community.

(5) Cost Realism and Project Management Plan: The overall estimated costs should be clearly justified and appropriate for the technical complexity of the effort. Evaluation will consider the value of the research to the government and the extent to which the proposed management plan will effectively achieve the capabilities proposed.

The Government reserves the right to select all, some, or none of the proposals received in response to this solicitation and to make awards without discussions with offerors; however, the Government reserves the right to conduct discussions if the Source Selection Authority later determines them to be necessary. Proposals identified for funding may result in a contract, grant, cooperative agreement, or other transaction depending upon the nature of the work proposed, the required degree of interaction between parties, and other factors. If warranted, portions of resulting awards may be segregated into pre-priced options.

The administrative addresses for this BAA are:

Fax: 703-741-7804 Addressed to: DARPA/IPTO, BAA 05-43
Electronic Mail: baa05-43@darpa.mil
Electronic File Retrieval: http://www.darpa.mil/ipto/Solicitations/solicitations.htm
Mail to: DARPA/IPTO
ATTN: BAA 05-43
3701 N. Fairfax Drive
Arlington, VA 22203-1714